

HVAC-Aware Occupancy Scheduling

BoonPing Lim, Menkes van den Briel, Sylvie Thiébaux,
Optimisation Research Group, NICTA
Research School of Computer Science, ANU
`first.last@anu.edu.au`

Scott Backhaus and Russell Bent
Defense Systems and Analysis Division
Los Alamos National Laboratory
`{backhaus, rbent}@lanl.gov`

Abstract

Energy consumption in commercial and educational buildings is impacted by group activities such as meetings, workshops, classes and exams, and can be reduced by scheduling these activities to take place at times and locations that are favorable from an energy standpoint. This paper improves on the effectiveness of energy-aware room-booking and occupancy scheduling approaches, by allowing the scheduling decisions to rely on an explicit model of the building's occupancy-based HVAC control. The core component of our approach is a mixed-integer linear programming (MILP) model which optimally solves the joint occupancy scheduling and occupancy-based HVAC control problem. To scale up to realistic problem sizes, we embed this MILP model into a large neighbourhood search (LNS). We obtain substantial energy reduction in comparison with occupancy-based HVAC control using arbitrary schedules or using schedules obtained by existing heuristic energy-aware scheduling approaches.

1 Introduction

Heating, ventilation and air-conditioning (HVAC) systems are responsible for 20% of the USA total energy consumption (Pérez-Lombard, Ortiz, and Pout 2008) and account for one hundred billion dollars/year electrical expenditure. These high energy costs together with rising environmental pollution levels call for innovative computational sustainability research focused on improving energy management in buildings.

Recent studies show that building HVAC consumption can be significantly reduced by adopting occupancy-based control strategies that exploit measured or predicted occupancy information (Erickson et al. 2009; Agarwal et al. 2010; West, Ward, and Wall 2014). For instance, model *predictive* control strategies determine supply air flow rate and temperature over longer time horizons so as to optimise energy consumption, whilst remaining within air flow and temperature bounds that reflect the predicted occupancy of various building zones (Goyal, Ingley, and Barooah 2013).

Another recent research line investigates the *proactive* control of *occupancy* in order to minimise HVAC consumption (Kwak et al. 2013; Majumdar, Albonese, and Bose 2012; Pan et al. 2012). Many office and university buildings offer

some scope for occupancy control via their room booking and scheduling systems. For instance meetings, exams, use of special purpose rooms, and other short-term activities can be scheduled to occur at times and in rooms that are beneficial from an energy standpoint. Unfortunately, existing occupancy scheduling approaches assume conventional HVAC control strategies (Kwak et al. 2013). Moreover, since reasoning using explicit models of the HVAC and the building is computationally expensive, they typically adopt suboptimal scheduling strategies guided by proxies for the optimisation criterion. One such proxy is the minimisation of the number of rooms used and of the time gap between successive meetings; it is used to guide the search towards solutions that take advantage of thermal inertia and schedule meetings to take place back-to-back in as few rooms as possible (Pan et al. 2012; Majumdar, Albonese, and Bose 2012).

This paper presents an approach which aims at combining the strengths of the above research directions. A naïve combination would be to schedule occupancy in a first phase using existing methods, and then control HVAC based on this occupancy schedule. In contrast, we model and solve the *joint* HVAC control and occupancy scheduling problem. This results in an integrated approach whose benefits exceed the naïve superposition of both of its parts, as the scheduling fully exploits the capabilities of the underlying occupancy-based HVAC control across available times and locations.

In more detail, the joint problem we consider is that of deciding the respective times and locations (rooms) of a set of meetings or similar activities, as well as the HVAC supply air temperature and air flow rate for each zone and time, in such a way as to optimise the overall HVAC consumption over a long time horizon. The schedule complies with the HVAC and building dynamics models, and with comfort and air-flow rate bounds that depend on the scheduled zone occupancy. It also satisfies typical meeting scheduling constraints, including constraints on meetings times, on meeting locations (e.g. room capacity, equipment availability), and on participant attendance conflicts. Once built, such a schedule can be used in any building equipped with an occupancy-based HVAC controller by simply complementing the occupancy forecast (Mamidi, Chang, and Maheswaran 2012) with the occupancy information captured in the schedule. Our approach can even be used with a conventional HVAC control system by constraining the bounds on supply air flow

rate/temperature and the temperature setpoints to be those found in the optimal solution to the joint problem.

To ensure the existence of a feasible control (for adequately sized HVACs) and improve on current occupancy-based HVAC control practices, we introduce a standby mode enabling the HVAC to re-activate at night if this is necessary to meet the temperature bounds of an early morning meeting or results in reduced consumption. To address the challenges caused by the presence of non-linear HVAC control constraints, we relax them in a principled way to obtain a mixed-integer linear (MILP) model guaranteed to provide a lower bound on the objective function. To tame the problem complexity further and scale to large problems, we combine the MILP model with Large Neighbourhood Search (LNS). LNS destroys parts of the schedule and MILP repairs them.

Our experiments illustrate the circumstances under which the standby mode is beneficial, demonstrate the approach’s scalability, and show the superiority (over 50% consumption reduction) of our joint MILP model compared to heuristic scheduling solutions and to more naïve integrations of meeting scheduling and occupancy-based HVAC control.

To summarise, the main contributions of the paper are: a) an efficient MILP model for occupancy-based HVAC control which can be used in a range of applications, b) a scalable and effective approach to occupancy scheduling which exploits the capabilities of HVAC control and c) experiments showing a substantial reduction in energy consumption over the state of the art.

2 Occupancy-Based HVAC Control

This section introduces the occupancy-based HVAC control aspects of our MILP model. Section 3 will cover the scheduling aspects. We describe in turn the type of HVAC system we focus on along with the objective function we consider, the bounds the control must comply with, the effect of the control on the building thermal dynamics, and the linear relaxation of the non-linear constraints.

Our HVAC control model builds on well-accepted work by Goyal et al. (2012; 2013), but introduces a number of significant changes required to be suitable as a subcomponent of our more complex joint scheduling and control model. Notably, the control model studied by (Goyal, Ingley, and Barooah 2013) is a purely continuous *non-linear* model which does not consider occupancy. Our experiments using the IPOPT solver (Wächter and Biegler 2006) revealed that this model is impractical for our application, in terms of both memory and run-time requirements. In contrast, our model is an efficient linear model that incorporates discrete variables to capture occupancy and our standby mode.

2.1 VAV System and Objective Function

Following Goyal et al. (2013), we consider variable-air-volume (VAV) based HVAC systems, which are widely deployed in commercial buildings. With such systems, the building is divided into a number of zones (or locations) $L \subseteq \mathbb{N}$, each of which can be an individual room or a group of rooms. To simplify notation, we assume that each zone corresponds to a single room. Figure 1 shows a schematic of a VAV system connected to two building zones.




Figure 1: VAV System with Two Zones

Let $K = \{1 \dots n\}$ be a finite set of discrete time steps considered over the optimisation horizon. For simplicity, we assume that successive time steps are separated by a fixed duration $\Delta t \in \mathbb{R}^+$; that is, $\forall k \in K$, we have $t_k \in \mathbb{R}^+$ and $t_k - t_{k-1} = \Delta t$. The objective is to minimise the total energy consumed over the optimisation horizon:

$$\text{minimise: } \sum_{k \in K} e_k \quad (1)$$

where e_k is the energy consumed at time step k :

$$e_k = p_k \times \Delta t \quad \forall k \in K \quad (2)$$

The power p_k is consumed by the three main operations shown in Figure 1 and detailed below: the air conditioning operation performed centrally by the air handling unit (AHU) consumes p_k^{Cond} , the fan operation, also performed centrally, consumes p_k^{Fan} , and the reheating operation performed locally at each zone $l \in L$ by the zone’s VAV unit consumes $p_{l,k}^{Heat}$ at each zone:

$$p_k = \left(p_k^{Cond} + p_k^{Fan} + \sum_{l \in L} p_{l,k}^{Heat} \right) \quad \forall k \in K \quad (3)$$

Air Conditioning Operation. The air handling unit (AHU) admits a mixture of outside air at temperature T_k^{OA} and return air, and conditions it to a pre-set conditioned air temperature T^{CA} (usually 12.8C). The conditioned air is then distributed through the supply duct to the VAV unit at each zone. The AHU consumption p_k^{Cond} is the power consumed in cooling the total air flow required. Let $a_{l,k}^{SA}$ denote the air flow rate required by location l at time step k and C^{pa} the heat capacity of air at constant pressure (1.005 kJ/kg·K):

$$p_k^{Cond} = C^{pa} (T_k^{OA} - T^{CA}) \sum_{l \in L} a_{l,k}^{SA} \quad \forall k \in K \quad (4)$$

Fan Operation. The supply fan, driven by a variable frequency drive, maintains a constant static pressure in the supply duct. When the opening of the VAV dampers increases to pull in more air flow into the conditioned space (resp. decreases to pull less air flow), the fan speeds up (resp. slows down). The fan consumption is the power consumed to push the total air flow required through the supply duct, which is proportional to the sum of the air flow rates $a_{l,k}^{SA}$ required over all locations. Let β be the fan coefficient (0.65):

$$p_k^{Fan} = \beta \sum_{l \in L} a_{l,k}^{SA} \quad \forall k \in K \quad (5)$$

Reheating Operation. Each zone l has a VAV unit connected to the supply duct. The unit is equipped with continuously adjustable valves and reheat coils. These enable regulating the air flow rate $a_{l,k}^{SA}$ into the zone and modulating the supply air temperature $T_{l,k}^{SA}$ to maintain the zone temperature within given bounds, if necessary by reheating the supply air. Here we consider the power $p_{l,k}^{Heat}$ consumed by the reheating process to heat the supply air from the conditioned temperature T^{CA} to an appropriate location supply air temperature $T_{l,k}^{SA}$.

$$p_{l,k}^{Heat} = C^{pa} (T_{l,k}^{SA} - T^{CA}) a_{l,k}^{SA} \quad \forall l \in L, k \in K \quad (6)$$

Decision Variables. The two key HVAC decision variables are the supply air flow rate $a_{l,k}^{SA}$ and temperature $T_{l,k}^{SA}$ at each location $l \in L$ and time step $k \in K$. We determine an optimal control for these variables, given occupancy information and bounds on supply air temperature, supply air flow rate, and room temperature during vacant and occupied periods. Below we will introduce a third decision variable ($w_{l,k}$) to decide when the HVAC should activate at night, which in turn will influence the bounds. When taking the HVAC control model in isolation, the building occupancy is an *input* to the model. When we integrate scheduling to the model in Section 3, occupancy will become a *variable*.

2.2 Temperature and Air Flow Bounds

We now model the constraints on the temperature, supply air temperature and supply air flow rate in each location, as a function of the location occupancy and the time of the day. We introduce the auxiliary variable $T_{l,k} \in \mathbb{R}$ representing the *actual* temperature in location $l \in L$ at time step $k \in K$, and the boolean input $z_{l,k}$ which is true iff l is occupied at time step k . When a location is not occupied, its temperature can lie freely within a wide temperature range $[T^{unocc,lb}, T^{unocc,ub}]$, whilst the temperature is otherwise constrained to lie within a more restricted comfort range $[T^{unocc,lb} + C^{lb}, T^{unocc,ub} - C^{ub}]$, where C^{lb} and C^{ub} are appropriate constants. This constraint is expressed as follows: $\forall l \in L, k \in K$

$$T^{unocc,lb} + C^{lb} z_{l,k} \leq T_{l,k} \leq T^{unocc,ub} - C^{ub} z_{l,k} \quad (7)$$

Further, the supply air temperature and flow rate at each location are constrained in a way that depends on the HVAC operating mode at the current time step. We have two operating modes: *active* and *standby*. Let $K^s \subseteq K$ be the set of time steps that fall within standard operating hours (6am to 6pm). During standard hours ($k \in K^s$) the HVAC is always in active mode. The supply air temperature $T_{l,k}^{SA}$ at location l must fall within $[T^{CA}, T^{SA,ub}]$. The supply air flow rate $a_{l,k}^{SA}$ must fall within $[a_l^{SA,lb}, a^{SA,ub}]$ where the upper bound is the air flow rate obtained when the dampers are fully open, and the lower bound is a constant (depending on the area size of the location and on the return air ratio) necessary to ensure that the minimal fresh outside air requirements are met. This yields the constraints:

$$T^{CA} \leq T_{l,k}^{SA} \leq T^{SA,ub} \quad \forall l \in L, k \in K^s \quad (8)$$

$$a_l^{SA,lb} \leq a_{l,k}^{SA} \leq a^{SA,ub} \quad \forall l \in L, k \in K^s \quad (9)$$




Figure 2: Thermal Model

Outside business hours ($k \in K \setminus K^s$), the HVAC is in stand-by mode and will only activate if this enables or lowers the cost of satisfying a future constraint. For instance, it could activate at night and benefit from the low outside night temperature to more cheaply cool the supply air to meet the temperature bounds in (7) for an early morning meeting. Note that this is different from conventional operations where HVACs are always off outside hours; as our experiments will show, the standby mode enables model-predictive approaches to occupancy-based control to meet constraints and save energy. The decision of whether or not HVAC activation is required by location l is represented by the boolean decision variable $w_{l,k}$. The presence of these boolean variables makes our model a mixed-integer model. When $w_{l,k}$ is true, the supply air flow rate and temperature are constrained to lie within $[T^{CA}, T^{SA,ub}]$ and $[a_l^{SA,lb}, a^{SA,ub}]$, respectively, and when $w_{l,k}$ is false, $a_{l,k}^{SA}$ is set to zero and the value of $T_{l,k}^{SA}$ is irrelevant (and for simplicity may as well also be zero). This is captured by the following constraints:

$$T^{CA} w_{l,k} \leq T_{l,k}^{SA} \leq T^{SA,ub} w_{l,k} \quad \forall l \in L, k \in K \setminus K^s \quad (10)$$

$$a_l^{SA,lb} w_{l,k} \leq a_{l,k}^{SA} \leq a^{SA,ub} w_{l,k} \quad \forall l \in L, k \in K \setminus K^s \quad (11)$$

2.3 Building Thermal Dynamics

Having defined the space of decisions as the supply air flow rate $a_{l,k}^{SA}$, the supply air temperature $T_{l,k}^{SA}$ and the HVAC activation requirement $w_{l,k}$ at each location and time step, we now model the impact of these decisions on the building thermal exchanges. To model the thermal dynamics of the building, we adopt a computationally efficient lumped RC-network (Gouda, Danaher, and Underwood 2000) which incorporates the thermal resistance and capacitance of each zone and between adjacent zones, as well as the solar gain and the internal heat gain in each zone – in particular the heat gain arising from occupancy. For the sake of simplicity, we ignore humidity and infiltration.

The principles behind the model are represented in Figure 2. Figure 2a shows the zone structure that we adopt. Zone l is separated by a wall and a window from zone $z1$ and by a wall from zones $z2$, $z3$, and $z4$, which could represent either indoor or outdoor zones. It is also separated by the ceiling and floor from zones c and f which are above and below zone l , respectively. Zone l has a capacitance C_l that models the heat capacity of the air in the

zone. It also has a solar gain $Q_{l,k}^s$ and heat gain $Q_{l,k}^p$ at time step k . Moreover, the inner (resp. outer) wall separating l from zone $z \in Z = \{z1, z2, z3, z4, f, c\}$ has a capacitance C_l^z (resp. C_l^l), resistance R_l^z (resp. R_l^l), and temperature $T_{l,k}^z$ (resp. $T_{z,k}^l$) at time step k . The window has a resistance R_l^w . Finally, the internal node between the inner and outer walls separating l from $z \in \{z1, z2, z3, z4\}$ has a constant resistance $R_l^{mid,z}$. Capacitances, resistances, solar gain, and (in this section) occupant heat gain are inputs to the model whilst temperatures are auxiliary variables. The interaction between zones is modeled using a lumped RC-network. Specifically, we use 3R2C for walls separating two zones, 2R1C for the ceiling and floor and 1R for windows. The lumped network for Figure 2a is given in Figure 2b.

The lumped network translates into a set of coupled difference equations which can be summarised as follows. The first difference equation defines the temperature $T_{l,k}$ in zone l at time step k as a function of the location, inner walls, ceiling, floor and outdoor temperatures at the previous time step, of the heat gain $Q_{l,k-1}^p$ at the previous time step and of the enthalpy $\Delta H_{l,k-1}$ of the location due to the supply air:

$$\begin{aligned} \frac{C_l}{\Delta t} (T_{l,k} - T_{l,k-1}) = & - \left[\sum_{z \in Z} \frac{1}{R_l^z} + \frac{1}{R_l^w} \right] T_{l,k-1} \\ & + \sum_{z \in Z} \frac{T_{l,k-1}^z}{R_l^z} + \frac{T_{k-1}^{OA}}{R_l^w} + Q_{l,k-1}^p + \Delta H_{l,k-1} \end{aligned} \quad (12)$$

The heat gain $Q_{l,k}^p$ is simply the heat gain q^p generated per person (75W), times the number of occupants ppi,k :

$$Q_{l,k}^p = q^p \times ppi,k \quad (13)$$

Ignoring humidity, the enthalpy is defined as follows:

$$\Delta H_{l,k} = C^{pa} a_{l,k}^{SA} (T_{l,k}^{SA} - T_{l,k}) \quad (14)$$

The remaining difference equations define the temperatures $T_{l,k}^z$ and $T_{z,k}^l$ of the inner and outer walls at time step k as a function of each other and of the location temperature $T_{l,k-1}$ at the previous time step. Taking $z = z1$ in the example of Figure 2a:

$$\begin{aligned} \frac{C_{z1}^l}{\Delta t} (T_{z1,k}^l - T_{z1,k-1}^l) = & - \left[\frac{1}{R_{z1}^l} + \frac{1}{R_l^{mid,z1}} \right] T_{z1,k-1}^l \\ & + \frac{T_{z1,k-1}}{R_{z1}^l} + \frac{T_{l,k-1}^{z1}}{R_l^{mid,z1}} + Q_{k-1}^s \end{aligned} \quad (15)$$

The definition of $T_{l,k}^{z1}$ is symmetrical except for the absence of solar gain Q_{k-1}^s . The equations for the other walls, and the ceiling and floor are similar, so we omit them here.

2.4 MILP Relaxation

Observe that the model as presented so far is a mixed-integer *non-linear* (MINLP) model. This is because of the bilinear terms $a_{l,k}^{SA} T_{l,k}^{SA}$ and $a_{l,k}^{SA} T_{l,k}$ in Equations 6 and 14. From a computational standpoint, it is better to relax these equations so as to obtain a MILP for which effective solvers exist that are guaranteed to return a lower bound on the globally optimal MINLP objective. To obtain a suitable MILP, we

use the linear programming relaxation of bilinear terms introduced by McCormick (1976). This relaxation introduces a new variable v for the bilinear term xy together with four inequalities that define its convex envelope using the bounds $[x, \bar{x}]$ and $[y, \bar{y}]$ on each of the two variables involved:

$$\begin{aligned} v & \geq \underline{xy} + \underline{yx} - \underline{xy} \\ v & \geq \bar{xy} + \bar{yx} - \bar{xy} \\ v & \leq \underline{xy} + \bar{yx} - \underline{xy} \\ v & \leq \bar{xy} + \underline{yx} - \bar{xy} \end{aligned}$$

Hence, our MILP model is obtained by replacing the bilinear terms $a_{l,k}^{SA} T_{l,k}^{SA}$ and $a_{l,k}^{SA} T_{l,k}$ in Equations 6 and 14 with new variables and adding the corresponding convex envelope definitions. The relevant bounds are:

- $a_{l,k}^{SA} \in [a_{l,k}^{SA}, \bar{a}_{l,k}^{SA}] = \begin{cases} [a^{SA,lb}, a^{SA,ub}] & \text{for } k \in K^s \\ [0, a^{SA,ub}] & \text{for } k \in K \setminus K^s \end{cases}$
- $T_{l,k}^{SA} \in [\underline{T}_{l,k}^{SA}, \bar{T}_{l,k}^{SA}] = \begin{cases} [T^{CA}, T^{SA,ub}] & \text{for } k \in K^s \\ [0, T^{SA,ub}] & \text{for } k \in K \setminus K^s \end{cases}$
- $T_{l,k} \in [\underline{T}_{l,k}, \bar{T}_{l,k}] = [T^{unocc,lb}, T^{unocc,ub}]$ for $k \in K$

This concludes the description of our MILP model for occupancy-based HVAC control. Given the occupancy ppi,k and $z_{l,k}$, and the external temperature T_k^{OA} , it controls the supply air flow rate $a_{l,k}^{SA}$ and temperature $T_{l,k}^{SA}$ and decides when a location requires HVAC activation $w_{l,k}$ out of the standby mode, in such a way as to optimise the total energy consumption $\sum_{k \in K} e_k$. The strengths of this model are its integration of realism and computational efficiency, its adequacy as a component of occupancy scheduling and other more complex models, and its optional ability to activate out of the standby mode when this improves consumption.

3 Occupancy Scheduling

Until now, zone occupancy over time was a model input. We now present our joint HVAC control and meeting scheduling model, in which occupancy is a decision variable.

Let $M \subseteq \mathbb{N}$ be a set of meetings to be scheduled to take place at the locations in L during the time horizon K . Each meeting $m \in M$ is characterised by the following inputs: its duration $\tau_m \in \mathbb{N}$ (number of time steps), the set of allowable time steps $K_m \subseteq K$ at which it can start, the set of allowable locations $L_m \subseteq M$ at which it can take place, and its set of attendees $P_m \subseteq A$, for some appropriate set of attendees A . In addition, let $N \subseteq 2^M$ be the set of meeting sets which have at least one attendee in common, that is $N = \{M_i \subseteq M \mid \forall m, m' \in M_i, P_m \cap P_{m'} \neq \emptyset\}$. In practice, only all pairs of incompatible meetings are needed. Note that the sets K_m and L_m can be used to encode a variety of situations, such as room capacity requirements and availability of special equipment such as video conferencing, as well as time deadlines for the meeting occurrence and attendee availability constraints.

The main scheduling variable is the boolean decision variable $x_{m,l,k}$ which is true iff meeting $m \in M$ is scheduled to take place at location $l \in L_m$ starting at time step $k \in K_m$. The scheduling part of the model interacts with the HVAC control part via the auxiliary variables $z_{l,k}$, which, as before,

is true iff location l is occupied at time step k , and $pp_{l,k} \in \mathbb{N}$, which, as before, represents the number of occupants at location l at time step k . These terms are used in Equations 7 and 13, respectively, but are now variables rather than inputs.

The set of MILP scheduling constraints are the following. The first constraint ensures that all meetings are scheduled to occur exactly once within the range of allowable locations and start times:

$$\sum_{l \in L_m, k \in K_m} x_{m,l,k} = 1 \quad \forall m \in M \quad (16)$$

The second constraint ensures that if a location is occupied by a meeting then it is exclusively occupied by this meeting during its entire duration:

$$\sum_{\substack{m \in M, k' \in K_m \\ \text{such that} \\ l \in L_m \text{ and } k - \tau_m + 1 \leq k' \leq k}} x_{m,l,k'} \leq z_{l,k} \quad \forall l \in L, k \in K \quad (17)$$

As a result, no two meetings can occupy the same location at the same time step. Observe that (17) also determines the occupancy variable $z_{l,k}$ used in the occupancy-based HVAC control part of the joint model.

The following constraint establishes the number of occupants $pp_{l,k}$ of each location l at each time step k :

$$\sum_{\substack{m \in M, k' \in K_m \\ \text{such that} \\ l \in L_m \text{ and } k - \tau_m + 1 \leq k' \leq k}} x_{m,l,k'} \times |P_m| = pp_{l,k} \quad \forall l \in L, k \in K \quad (18)$$

This is used in equation 13 to establish the internal heat gain arising from occupancy.

Finally, the last constraint ensures that meetings with an intersecting attendee set cannot overlap in time:

$$\sum_{\substack{m \in \nu, l \in L_m, k' \in K_m \\ \text{such that} \\ k - \tau_m + 1 \leq k' \leq k}} x_{m,l,k'} \leq 1 \quad \forall k \in K, \nu \in N \quad (19)$$

Our joint HVAC control and occupancy scheduling model is simply obtained by adding equations 16-19 the HVAC control model given by equations 1-15 (with 6 and 14 linearised). The model optimises the total energy consumed not only over the HVAC decision variables $a_{l,k}^{SA}$, $T_{l,k}^{SA}$ and $w_{l,k}$ as before, but also over the scheduling decision variables $x_{m,l,k}$. A building occupancy-based HVAC controller need only use the schedules $x_{m,l,k}$ produced. A conventional controller may instead use the bounds on room temperature and supply air flow rate/temperature determined by equations 7, 10 and 11 as setpoints.

This concludes the description of our joint HVAC control and occupancy scheduling model. The next section experimentally investigates its benefits in terms of energy reduction, in comparison with more naïve integrations of scheduling and HVAC control. Section 5 shows that by embedding this model into a large neighbourhood search, we can scale to timetabling problems of practical relevance.

4 Benefits of the Model

Our experiments aim at explaining the usefulness of the standby mode and at demonstrating that our HVAC-aware scheduling model leads to significant consumption reduction (50% to 70% in our experiments) when compared to




Figure 3: HVAC control with and without standby mode.

occupancy-based HVAC control using arbitrary schedules or energy-aware schedules generated by heuristic methods. Experiments are conducted over 5 summer days with a row of 4 co-located zones, each consisting of a single 60 m² room with a capacity of 30 people. The zones differ by a high or low value for their thermal resistance and capacitance. The two end zones have three outside walls and the middle two zones have two. The duration between successive time steps is $\Delta t = 30\text{min}$, giving more than enough time for thermal effects to occur. Shorter durations did not significantly affect the results. The input data and parameters are available from the first author. The MILP models are solved using Gurobi 5.6 (2014). All experiments were conducted on a cluster consisting of 2× AMD 6-Core Opteron 4184, 2.8 GHz with 64 GB of memory.

4.1 Usefulness of Standby Mode

We start by illustrating the usefulness of the standby mode. In conventional operations, HVACs are usually switched on a few hours prior to start of business (6am) and are turned off in the evening (6pm) and at night. Model predictive control strategies are capable of pre-cooling a zone, but only when the HVAC is switched on. Our standby mode enables the HVAC to self activate outside business hours to provide additional pre-cooling when this is beneficial. Because HVAC consumption is highly dependent on the temperature gap between the outdoor temperature and the conditioned air temperature, pre-cooling at night, when the outdoor air temperature is cooler, can reduce energy consumption. The following experiment shows that such pre-cooling can be beneficial not only for early morning meetings, but also, more surprisingly, for late afternoon meetings.

Figure 3 compares the operations of the HVAC optimally controlled by the model in Section 2 with standby mode (S) and without standby mode (N). For this experiment, a single meeting is scheduled to occur between 16:00-17:00 in a given zone on a given day. Observe that when the HVAC is running with the standby mode enabled, it activates as early as 02:30 and pushes between 2.2 and 1.2 kg/s of supply air at 12.8 C to bring down the zone temperature to approximately

19°C by 09:00. Between 02:30 and 6:00, the outdoor temperature lies between 15 and 17°C, which is about 2-4°C higher than the 12.8°C conditioned air temperature. Without the standby mode, supply air is pushed into the room at a higher average rate between 2.0 and 1.5 kg/s right after the HVAC is turned on at 06:00, which, as the outdoor temperature is higher at that time (18-22°C), requires a higher rate of energy consumption. During the day, the zone temperature increases slightly due to the daytime thermal gain, and at 15:00, one hour before the meeting starts, the room is pre-cooled again. This time, the standby-mode enabled HVAC requires cooling about half the amount of supply air, which brings significant energy savings since the outside temperature is around 36°C. Altogether, the standby mode reduces consumption by 11.9% (12kWh) on this example.

As shown above, a standby-mode-enabled HVAC can be effective in areas with high diurnal temperature variation. In addition to decreasing energy consumption, it can provide pre-cooling at off-peak electricity cost. For organisations that are charged by electricity suppliers according to their peak consumption, another benefit of the standby mode is that it can help smooth the peak that is regularly observed at the start of the operating hours.

4.2 Joint Model vs Simpler Models

Whilst the standby mode is beneficial, the much larger gains in our approach stem from the joint model: we now compare our joint model with simpler approaches representative of the existing literature on occupancy-based HVAC control and energy-aware meeting scheduling, and observe a 50%-70% energy consumption improvement. Specifically, we consider a set of timetabling problems derived from the PATAT (2002) Melbourne University dataset and compare the optimal (O) solutions produced by the joint model in Section 3, with those produced by giving arbitrary (A) schedules and heuristic (H) energy-aware schedules as input to the HVAC control model in Section 2. Several authors have argued that scheduling meetings back to back in as few rooms as possible is a suitable heuristic that takes advantage of thermal inertia to reduce energy consumption (Kwak et al. 2013; Majumdar, Albonese, and Bose 2012; Pan et al. 2012). In line with this, the heuristic we compare to minimise the number of rooms used and the time gap between meetings in these rooms, subject to the scheduling constraints 16-19.¹ In all three cases (A,H,O), we run the HVAC control model with standby mode (S) and without it (N), resulting in six different methods labeled AN, AS, HN, HS, ON, OS, where for example, HS denotes HVAC control with standby mode using heuristic schedules.

To examine problems with different degree of constrainedness, we extracted 70 problem instances from the PATAT dataset, consisting of 40 instances of 10 meetings each, 20 instances of 20 meetings each, and 10 instances of 50 meet-

¹Majumdar et al. (Majumdar, Albonese, and Bose 2012) observe that the single most important predictor of performance is good match between the room capacity and the size of the meetings. However this does not play a role in our experiments since all four rooms have the same capacity.

Strategy	Average energy consumption (kWh)	Excess consumption vs. baseline
AN	212.14	74.84%
AS	199.94	64.78%
HN	184.26	51.86%
HS	177.32	46.14%
ON	124.13	2.30%
OS	121.34	baseline

Table 1: Comparison of arbitrary (A), heuristic (H), and optimal (O) scheduling strategies over HVAC with (S) and without (N) standby mode.

ings each. All meetings have up to 30 attendees, a 1.5h duration and an allowable time range of one or two random days (09:00-17:00) within the 5 days of the experiment.

The AN/AS results are obtained by selecting, for each instance, an arbitrary schedule consistent with the scheduling constraints 16-19 in Section 3 and using it as input to the occupancy-based HVAC control model in Section 2. Similarly, the HN/HS results are obtained by selecting the schedule optimising the heuristic among those consistent with the scheduling constraints, and using it as input to the occupancy-based HVAC control model. The ON/OS results are obtained by solving the joint model for each instance.

Table 1 shows, for each of the 6 approaches, the average energy consumption per room over the 70 instances, and the percentage excess consumption taking OS as the baseline. The results show a clear improvement as we move from arbitrary schedules (AN/AS), that are currently the norm with room booking systems, to energy aware schedules (HN/HS), and a much greater improvement when these schedules take into account the capabilities of occupancy-based HVAC control (ON/OS). The interactions between the various scheduling constraints, the thermal dynamics of the building and the HVAC control are so complex that heuristic methods can only achieve a fraction of the performance of the global optimisation methods enabled by our MILP model. As expected, the gain conferred by the standby mode decreases as we move to schedules that make better time and location decisions. Similarly, we observed that for more constrained problems (e.g. with 50 meetings), the standby mode is more effective, because there is a greater likelihood that meetings need to be scheduled in rooms that require higher cooling load which the standby mode can mitigate by pre-cooling.

5 Scaling to Large Problems

MILP enables us to easily manage the tightly constrained interactions between meeting scheduling and its direct impact on energy consumption. However, it only allows us to solve small problem instances in reasonable time. To scale to problem sizes that universities face when scheduling exams, we developed a hybrid solution that embeds the MILP model into a large neighbourhood search (LNS) (Shaw 1998).

LNS is a local search metaheuristic, which iteratively improves an initial solution by alternating between a destroy step and a repair step. The main idea behind LNS is that a large neighbourhood allows the heuristic to easily navigate




Figure 4: Performance of LNS, MILP and HS.

through the solution space and escape local minima even when the problem is highly-constrained. One important decision when implementing the destroy step is determining the amount of destruction. If too little is destroyed the effect of a large neighbourhood is lost and if too much is destroyed then the approach turns into repeated re-optimization. Another important decision is whether the repair step should be optimal or not. An optimal repair will be slower than a heuristic, but may potentially lead to high quality solutions in a few iterations. As a result, some parameter tuning will be essential in achieving good performance overall.

In our destroy step, we remove all meetings in two, three, or four randomly selected zones. This forms a subproblem that the repair step can effectively solve using MILP. We do, however, limit the MILP runtime to avoid excessive search during repair. That means we do not necessarily solve the subproblem to optimality, but given that MILP solvers are anytime algorithms, we do improve solution quality in many of the LNS iterations. We used the sequential model-based algorithm configuration (SMAC) methodology (Hutter, Hoos, and Leyton-Brown 2011) on an independent set of problems to optimise the parameters of the probability of the number of rooms to destroy and the MILP run time.

Figure 4 compares the average energy consumption obtained by LNS, MILP and the HS heuristic on 100 runs for each of 80 larger instances extracted from the PATAT dataset. These consist of 8 groups of 10 instances each, ranging from 20 to 500 1-1.5h meetings to be scheduled in 20 to 50 rooms over the 5 days. For each run, both MILP and LNS were seeded with HS as the initial solution and were given the same run-time limit of 15 minutes. The percentages in the figure show the average excess consumption of MILP and HS, taking LNS as the baseline. The bottom bars give the average excess over all instances and runs. The figure shows that LNS is capable of returning significantly better solutions on large problems. The performance gap increases as the problem scales, and is higher for problems that are neither too weakly nor too tightly constrained (e.g. 100M/20R, 200M/50R). For the largest in-

stances (500M/50R) MILP could not improve at all over the heuristic. MILP’s performance did not measurably improve with a 2h run-time.

6 Conclusion, Related and Future Work

In this paper we focus on meeting scheduling, but our model is more broadly applicable to scheduling occupant activities within specified time windows, and ultimately, can be integrated into a range of room booking and scheduling systems. To bring awareness of the capabilities of the building’s HVAC system to the scheduling process, we solve a joint HVAC control and occupancy scheduling problem. This problem involves determining the times and locations of a set of meetings, as well as the supply air temperature and air flow rate for each building zone, so as to minimize HVAC consumption. Existing approaches solve the HVAC control problem and the occupancy scheduling problem in isolation. While the joint problem is more challenging, it does achieve a much higher rate of energy savings. By combining LNS and MILP we are able to generate good solutions to large instances within 15 minutes, making the approach practical for university timetabling applications.

Combining constraint-based methods with LNS is not new. For example, Di Gaspero et al. (2013) used LNS with CP in balancing bike sharing systems and Le Bras et al. (2013) used LNS with MILP in planning for wildlife conservation. We are, however, unaware of its application in the smart buildings space.

Previous work on occupancy-based HVAC control treats occupancy information as an input parameter and not as a control variable (Agarwal et al. 2010; Brooks and Baroah 2014; Mady et al. 2011; Parisio et al. 2013; Xu et al. 2009). Our work is different as it incorporates both HVAC control and occupancy scheduling into a unified model. Moreover, our standby mode improves the feasibility and solution quality of model-predictive HVAC control methods.

Energy-aware scheduling methods typically take advantage of thermal inertia using heuristics such as minimizing the number of rooms used and/or assigning lower costs to meetings scheduled back to back in the same room (Pan et al. 2012; Kwak et al. 2013; Majumdar, Albonesi, and Bose 2012). An important limit of these works is that they all share a black-box modeling approach for calculating HVAC energy consumption. This black-box approach confines the search to a rather limited space that does not exploit the HVAC’s full capabilities. Another limitation is the assumption of an anonymous list of meeting participants, which leads to ignore the existence of meeting conflicts and the fundamental need for resolving them.

Directions for future work include: (1) extending our model to include more HVAC parameters such as humidity, (2) incorporating dynamic temperature bounds that adjust thermal comfort bounds based on outdoor temperature, (3) investigating the application of multi-objective optimisation to over-constrained problems where, additionally, violation of scheduling constraints needs to be minimised, and (4) integrating online scheduling and cancellation of meeting requests in real-time as in (Kwak et al. 2013).

Acknowledgments

Thanks to Pascal Van Hentenryck for pointing us to optimisation problems in the Smart Buildings space. Thanks to Pascal, Alban Grastien and the referees for useful comments. This work is supported by NICTA's Optimisation Research Group as part of the Future Energy Systems project. NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

- Agarwal, Y.; Balaji, B.; Gupta, R.; Lyles, J.; Wei, M.; and Weng, T. 2010. Occupancy-driven energy management for smart building automation. In *Proc. ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys)*, 1–6.
- Brooks, J., and Barooah, P. 2014. Energy-efficient control of under-actuated HVAC zones in buildings. In *Proc. American Control Conference (ACC)*, 424–429.
- Erickson, V. L.; Lin, Y.; Kamthe, A.; Brahme, R.; Surana, A.; Cerpa, A. E.; Sohn, M. D.; and Narayanan, S. 2009. Energy efficient building environment control strategies using real-time occupancy measurements. In *Proc. ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys)*, 19–24.
- Gaspero, L. D.; Rendl, A.; and Urli, T. 2013. Constraint-based approaches for balancing bike sharing systems. In *Proc. International Conference on the Principles and Practice of Constraint Programming (CP)*, 758–773.
- Gouda, M.; Danaher, S.; and Underwood, C. 2000. Low-order model for the simulation of a building and its heating system. *Building Services Engineering Research and Technology* 21(3):199–208.
- Goyal, S., and Barooah, P. 2012. A method for model-reduction of non-linear thermal dynamics of multi-zone buildings. *Energy and Buildings* 47:332–340.
- Goyal, S.; Ingle, H. A.; and Barooah, P. 2013. Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance. *Applied Energy* 106:209–221.
- Gurobi Optimization, I. 2014. Gurobi optimizer reference manual. <http://www.gurobi.com>.
- Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2011. Sequential model-based optimization for general algorithm configuration. In *Proc. International Conference on Learning and Intelligent Optimization (LION)*, 507–523.
- Kwak, J.-y.; Varakantham, P.; Maheswaran, R.; Chang, Y.-H.; Tambe, M.; Becerik-Gerber, B.; and Wood, W. 2013. Tesla: An energy-saving agent that leverages schedule flexibility. In *Proc. International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 965–972.
- LeBras, R.; Dilkina, B. N.; Xue, Y.; Gomes, C. P.; McKelvey, K. S.; Schwartz, M. K.; and Montgomery, C. A. 2013. Robust network design for multispecies conservation. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, 1305–1312.
- Mady, A. E.-D.; Provan, G.; Ryan, C.; and Brown, K. N. 2011. Stochastic model predictive controller for the integration of building use and temperature regulation. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, 1371–1376.
- Majumdar, A.; Albonese, D. H.; and Bose, P. 2012. Energy-aware meeting scheduling algorithms for smart buildings. In *Proc. ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys)*, 161–168. ACM.
- Mamidi, S.; Chang, Y.; and Maheswaran, R. T. 2012. Improving building energy efficiency with a network of sensing, learning and prediction agents. In *Proc. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 45–52.
- McCormick, G. P. 1976. Computability of global solutions to factorable nonconvex programs: Part I- convex underestimating problems. *Mathematical programming* 10(1):147–175.
- Melbourne University. 2002. PATAT 2002 Dataset. <http://www.or.ms.unimelb.edu.au/timetabling/>.
- Pan, D.; Yuan, Y.; Wang, D.; Xu, X.; Peng, Y.; Peng, X.; and Wan, P.-J. 2012. Thermal inertia: Towards an energy conservation room management system. In *Proc. IEEE International Conference on Computer Communications (INFOCOM)*, 2606–2610.
- Parisio, A.; Varagnolo, D.; Risberg, D.; Pattarello, G.; Molinari, M.; and Johansson, K. H. 2013. Randomized model predictive control for HVAC systems. In *Proc. ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys)*, 1–8.
- Pérez-Lombard, L.; Ortiz, J.; and Pout, C. 2008. A review on buildings energy consumption information. *Energy and Buildings* 40:394–398.
- Shaw, P. 1998. Using constraint programming and local search methods to solve vehicle routing problems. In *Proc. International Conference on Principles and Practice of Constraint Programming (CP)*, 417–431.
- Wächter, A., and Biegler, L. T. 2006. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106(1).
- West, S. R.; Ward, J. K.; and Wall, J. 2014. Trial results from a model predictive control and optimisation system for commercial building HVAC. *Energy and Buildings* 72:271–279.
- Xu, X.; Wang, S.; Sun, Z.; and Xiao, F. 2009. A model-based optimal ventilation control strategy of multi-zone vav air-conditioning systems. *Applied Thermal Engineering* 29(1):91–104.